

3-15-00

A

3-14-00

10544 U.S. PTO

Please type a plus sign (+) inside this box → ☐

PTO/IS/05 (4/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No. **MS-81 (116629.1)**
First Inventor or Application Identifier **R. VENKATESAN**
Title **A. TECHNIQUE FOR PRODUCING, THROUGH WATERMARKING, HIGHLY TAMPER-RESISTANT...**
Express Mail Label No. **EL44371403605**

etc.

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO: Assistant Commissioner for Patents
Box Patent Application
Washington, DC 20231

- ☒ Fee Transmittal Form (e.g., PTO/IS/17) (2 pps.)
(Submit an original and a duplicate for fee processing)
- ☒ Specification [Total Pages **60**]
(preferred arrangement set forth below)
 - Descriptive title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claim(s)
 - Abstract of the Disclosure
- ☒ Drawings (35 U.S.C. 113) [Total Sheets **8**]
(Formal)
- Oath or Declaration [Total Pages **1**]
 - ☐ Newly executed (original or copy)
 - ☐ Copy from a prior application (37 C.F.R. § 1.63(d))
(for continuation/divisional with Box 16 completed)
 - ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

- ☐ Microfiche Computer Program (Appendix) (if applicable, all necessary)
 - ☐ Computer Readable Copy
 - ☐ Paper Copy (identical to computer copy)
 - ☐ Statement verifying identity of above copies
- ☐ Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary)

ACCOMPANYING APPLICATION PARTS

- ☐ Assignment Papers (cover sheet & document(s))
- ☐ 37 C.F.R. § 3.73(b) Statement of Power of Attorney (when there is an assignee)
- ☐ English Translation Document (if applicable)
- ☐ Information Disclosure Statement (IDS)/PTO-1449
- ☐ Copies of IDS Citations
- ☐ Preliminary Amendment
- ☒ Return Receipt Postcard (MPEP 503) (2)
(Should be specifically itemized)
- ☐ Small Entity Statement(s) Status still proper and desired (PTO/IS/09-12)
- ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
- ☒ Other: Cover Letter (2 pps.)

*NOTE FOR ITEMS 1 & 12: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED BY 37 C.F.R. § 1.57, EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON BY 37 C.F.R. § 1.58.

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment.

- ☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No. _____ / _____
Prior application information: Examiner _____ Group / Art Unit: _____
For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS

☒ Customer Number or Bar Code Label **007265** or ☐ Correspondence address below
(Insert Customer No. or Attach bar code label here)

Name _____
Address _____
City _____ State _____ Zip Code _____
Country _____ Telephone _____ Fax _____

Name (Print/Type) **Pever L. Michaelson** Registration No. (Attorney/Agent) **30,090**
Signature *Pever L. Michaelson* Date **3/14/00**

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS SEND TO Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

10544 U.S. PTO
03/14/00

MICHAELSON & WALLACE
A Partnership of Intellectual Property and Technology Lawyers

Doc. No. MS-81(116629.1)

March 14, 2000

PATENT APPLICATION

Peter L. Michaelson
NJ, NY, PA Bars

Robert M. Wallace
NY, CA Bars

Christopher R. Balzan
CA Bar

Eric J. Agaard
CA Bar

Arien C. T. Ferrell
CA Bar

Jeremiah G. Murray
Patent Agent

Of Counsel:

John T. Peoples
NJ Bar

Ronald L. Drumheller
NY Bar

Edward M. Fink
NY, D.C. Bars

Assistant Commissioner of Patents
Box **Patent Application**
Washington, D. C. 20231

Sir:

Enclosed herewith for filing is the following
utility patent application:

Applicants: **Ramarathnam VENKATESAN, Vijay VAZIRANI**

Title of application: **A TECHNIQUE FOR PRODUCING, THROUGH
WATERMARKING, HIGHLY TAMPER-RESISTANT
EXECUTABLE CODE AND RESULTING
"WATERMARKED" CODE SO FORMED**

Total pages of specification: 60 (45 pages of specification,
14 pages of claims [1-50]
and 1 abstract page)

Sheets of drawing: 8 (FIGs. 1-10)

Docket No.: **MS-81(116629.1)**

Internet Address:
mail@mandw.com

Web Site:
http://www.mandw.com

Please reply:

- ☒ Parkway 109 Office Ctr.
328 Newman Springs Rd.
P. O. Box 8489
Red Bank, NJ 07701
Tel: 732-530-6671
Fax: 732-530-6584 or
732-530-6585
Videoconference:
732-224-0132 (ISDN)

- ☐ 1001 Partridge Drive
Suite 330
Ventura, CA 93003
Tel: 805-644-4035
Fax: 805-644-8046

PATENT APPLICATION FILING FEE CALCULATION

	<u>No. Filed</u>	<u>Less</u>	<u>Rate/Claim</u>	<u>Fee</u>
Total				
Claims	<u>50</u>	-20	30 x \$18.00	\$ 540.00
Independent				
Claims	<u>2</u>	-3	0 x \$78.00	\$ 00.00
Minimum Filing Fee				\$ 690.00
Multiple Dependency Fee (if applicable - \$260.00)				\$ 00.00

50% Reduction for Small Entity
(Independent Inventor, Non-profit
Corporation, or Small Business
Concern) - appropriate

verified statement attached	\$- 00.00
TOTAL FILING FEE	\$ 1,230.00
TOTAL FILING FEE ENCLOSED	\$ 00.00

Also enclosed herewith for filing in connection
with the enclosed application are:

XX Postcard Receipt (2); and
XX Utility Patent Application Transmittal.

Respectfully submitted,

MICHAELSON & WALLACE



Peter L. Michaelson, Attorney
Reg. No. 30,090
Customer No. 007265

EXPRESS MAIL CERTIFICATION

"Express Mail" mailing label number: EL443714036US
Date of deposit: March 14, 2000

I hereby certify that this paper or fee is being
deposited with the United States Postal Service "Express Mail
Post Office to Addressee" service under 37 CFR 1.10 on the
date indicated above and is addressed to the Assistant
Commissioner of Patents, **Box Patent Application**, Washington,
D.C. 20231


Signature of person making certification

Peter L. Michaelson
Name of person making certification

(MS81PTOLTR/139:ca)

-1-

**A TECHNIQUE FOR PRODUCING, THROUGH WATERMARKING, HIGHLY
TAMPER-RESISTANT EXECUTABLE CODE AND RESULTING
"WATERMARKED" CODE SO FORMED**

BACKGROUND OF THE DISCLOSURE

1. Field of the Invention

5 The invention relates to a technique including
both apparatus and an accompanying method, for forming
and embedding a hidden highly tamper-resistant
cryptographic identifier, i.e., a watermark, within
10 non-marked computer executable code, e.g., an application
program, to generate a "watermarked" version of that
code. This technique can also be used to tightly
integrate, in a highly tamper-resistant manner, other
pre-defined executable code, such as security code, as
15 part of the watermark, into the non-marked code in order
to form the watermarked code.

2. Description of the Prior Art

20 Over the past decade or so, personal computers
(PCs) have become rather ubiquitous with PC hardware and
software sales experiencing significant growth. However,
coincident with an ever widening market for PCs,
unauthorized copying of PC software, whether it be
25 application programs or operating systems, continues to
expand to rather significant proportions. Given that in
certain countries sales lost to such copying can
significantly exceed legitimate sales, over the years

software manufacturers have attempted to drastically reduce the incidence of unauthorized copying though, practically speaking, with only limited success.

5 One such technique, probably one of the oldest techniques used and usually rather ineffective, is simply to append a copyright and other legal proprietary rights notices to object code as distributed on mass (magnetic or optical) media. The intention in doing so is to deter
10 unauthorized copying by simply placing a third party on notice that a copy of the program, embodied by that code, is legally protected and that its owner may take legal action to enforce its rights in the program against that party to prevent such copying. These notices can be
15 readily discovered in program code listings and simply excised by the third party prior to copying and distributing illicit copies. Other such notices can be excised by a third party adversary from the software media itself and the program packaging as well. Though
20 these notices are often necessary in many jurisdictions to secure full legal remedies against third parties, in practice, these notices have provided little, if any, real protection against third party copying.

25 Another technique that is recently seeing increasing use is to require a PC, on which the program is to execute, to hold a valid digital "certificate" provided by the manufacturer of the program. The certificate will typically be loaded as a separate step
30 during manufacture of the PC. During initialization, the program will test the certificate and confirm its

authenticity and validity. If the certificate is authentic and valid, the program will continue to execute; otherwise, the program will simply terminate. Unfortunately, the certificate and associated testing routines are often very loosely bound to the remainder of the program code. Currently available software analysis tools can display execution flow among program instructions in a program under test. Consequently, with such tools, a programmer, with knowledge of an operational sequence implemented by the program and by analyzing a flow pattern inherent in that program, as it executes, can readily discern the program code that implements a certificate testing function. Once this code is detected, the programmer can readily excise that portion from the program itself and simply modify the remaining program code, by, e.g., by inclusion of appropriate jump instruction(s), to compensate for the excised portion; thus, totally frustrating the protection which the certificate was intended to provide against unauthorized copying. Once having done so, a third party adversary can then produce and distribute unauthorized, but fully executable, copies of the program free of all such protection. Thus, in practice, this approach has proven to be easily compromised and hence afforded very little, if any, real protection against illicit copying.

Other techniques have relied on using serialized hardware or other hardware centric arrangements to limit access to a program to one or more users at one particular PC and preclude that program from being loaded onto another PC. Generally, these

techniques, often referred to as "copy protect" schemes and which were popular several years ago, relied on inserting a writeable program distribution diskette, such as a floppy diskette, into a PC and then, during execution of an installation process from that diskette, have that PC store machine specific data, such as a hardware identification code, onto the diskette. Thereafter, during each subsequent installation of the program, an installation process would check the stored machine specific data on the installation diskette against that for a specific PC on which the program was then being installed. If the two pieces of data matched, installation would proceed; otherwise, it would prematurely terminate. Unfortunately, such schemes, while generally effective against unauthorized copying, often precluded legitimate archival copying as well as a legitimate installation of the program on a different PC. In view of substantial inconveniences imposed on the user community, such "copy protect" schemes quickly fell into disuse and hence were basically abandoned shortly after they first saw widespread use. Moreover, any such technique that relies on storing information on the distribution media itself during program installation is no longer practical when today software is distributed on massive read-only optical media, such as CDROM or, soon, digital video disk (DVD).

Therefore, given the drawbacks associated with copy protect and certificate based schemes, one would think that embedding an identifier of some sort into a program, during its manufacture and/or installation and

subsequently testing for that identifier during subsequent execution of an installed version of that program at a user PC, would hold promise.

5 However, for such an identifier based approach to be feasible, a need exists in the art for an identifier, such as a watermark, that can be tightly integrated into a program itself such that the watermark would be extremely difficult, if not effectively
10 impossible, for a third party to discern, such as through flow analysis, and then circumvent, such as by removal.

 In particular, such a watermark could be embedded in some fashion into a non-marked program.
15 Then, subsequently, at runtime of an installed version of that program at a user PC, a "secret" key(s) based cryptographic process could be used to reveal the presence of and test the watermark. The key(s) would be separately stored down, to the PC, as a software
20 value(s). If the correct watermark were then detected, execution of the installed program would continue; else, execution would halt. Fortunately, such an approach would likely impose essentially no burden on, and preferably be totally transparent to, the user, and not
25 frustrate legitimate copying.

 If such an identifier could be made sufficiently impervious to third party detection and tampering, then advantageously its use, with, for
30 example, such an approach, may well prove effective, in practice, at reducing unauthorized third party copying.

SUMMARY OF THE INVENTION

Our present invention advantageously satisfies this need and overcomes the deficiencies in the art through a watermark, containing, e.g., a relatively large number of executable routines, that is tightly integrated into a flow pattern of non-marked executable code, e.g., an application program, through randomly establishing additional control (execution) flows in the executable code and inserting a selected one of the routines along each such flow. Since a resulting flow pattern of the watermark is highly intertwined with (tightly spliced into) the flow pattern of the non-marked code, the watermark is effectively impossible to either remove from the code and/or circumvent. Furthermore, the code for the routines themselves is added in such a manner that the flow pattern of resulting "watermarked" code is not substantially different from that of the non-marked code. Hence, the watermark is also extremely difficult for a third party adversary to discern using, e.g., standard flow analysis tools and human inspection.

Advantageously, to enhance tamper-resistance of the watermarked code, each routine, that constitutes a portion of the watermark, can provide a pre-defined function such that, if that routine were to be removed from the marked code by, e.g., the third party adversary, then the marked code will prematurely terminate its execution.

In accordance with our specific inventive teachings, unmarked executable code which forms, e.g., an application program that is to be watermarked is first converted, using a conventional software flow analysis tool, into its corresponding flow graph. Predefined security code, typically constituting specific predetermined executable software code, is also converted, through use of the same tool, into its corresponding flow graph. The security code can itself constitute, for example: specific "watermark" code, i.e., executable code having as its primary, if not sole, purpose to form a portion of a watermark (i.e., distinct from the application program itself); a complete image of the entire application program itself; or just a portion of that program. In that regard, the unmarked executable code and the security code can each be formed of a different half of a common application program.

Thereafter, each of the flow graphs is k-partitioned to yield cluster flow graphs G' and H' , respectively (where k is a pre-defined integer, such as illustratively 1000 for a large application program then being watermarked), each having k clusters of nodes (each being a partition). M edges (links) (where M is typically a large pre-defined integer, such as illustratively 500,000 for that application program) are collectively inserted between corresponding pairs of randomly selected nodes in: (a) graphs G' and H' ; and, where desired, (b) different clusters solely in graph G' , and/or (c) different clusters solely in graph H' . For each edge, a routine is selected from a pre-defined

library, based on, e.g., minimizing adverse affects on
program flow, and its designation is inserted along that
edge in the flow graph(s), specifically at one of the
nodes associated with that edge. All the edges
5 collectively and effectively splice clustered flow
graphs G' and H' together into a single combined flow
graph. Executable code is then produced which
corresponds to that depicted in the single combined flow
graph. The watermark is collectively defined by the
10 routines and edges that have been inserted into the
unmarked code.

One illustrative heuristic for selecting each
specific pair of nodes, in, e.g., cluster graphs G' and
15 H' that are to be joined by an edge, is as follows.
First, randomly pick a node, U , in graph G' . With λ
being pre-defined as equaling (a number of edges that are
to transit between G' and H')/(a number of edges
connected to U), then, with a probability of $1-\lambda$,
20 randomly choose a node, Y , in graph H' . Then, with a
probability of λ , randomly choose a node, Z other than U ,
in graph G' . Finally, provide, as output, designations
of nodes Y and Z as a nodal pair.

25 During subsequent edge insertion, connect the
nodes for that edge together, e.g., nodes Y and Z , so as
to insert an edge extending between cluster flow
graphs G' and H' . Based on proper program flow, insert
an appropriate routine from the library along that edge
30 and at an appropriate node, in a graph, for that pair.

Repeat these node selection and insertion steps until all M edges and designations for associated routines are collectively added to cluster graphs G' and H' so as to fully splice both graphs and the associated routines into a single combined flow graph. Parameters k , M and λ are preferably kept in secret.

Each of these routines is predetermined, usually quite compact, requires relatively little execution time and executes a pre-defined, often self-contained operation, such as, e.g., computing a cryptographic key for use in printing or decoding a variable, or decrypting a ciphered variable. Each of the operations is chosen so as not to require much processing time; thus, not noticeably degrading execution of the watermarked program. Collectively, the routines that are inserted are such that, for proper execution of the watermarked program, they must all be executed and, to a certain extent, in a given sequence. In that regard, if any one or more of these routines is removed from the watermarked program, such as by a third party adversary, that program will gracefully terminate its execution.

To further frustrate its detection, the code for all the inserted routines is collectively scattered approximately uniformly throughout the "watermarked" program as that program is being constructed from its combined flow graph. In this manner, the routines will not be centralized in any one portion of the watermarked program. Furthermore, each of these routines is written

with standard code "obfuscation" techniques to further camouflage their functionality.

Advantageously, as a feature, the present invention can securely watermark any executable code, whether it forms, e.g., an application program, an operating system (O/S) or a software module.

BRIEF DESCRIPTION OF THE DRAWINGS

The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 depicts a simplified high-level block diagram of conventional watermarking process 100;

FIG. 2 depicts a simplified and illustrative flow graph produced through conventional program flow analysis;

FIG. 3 depicts a simplified high-level block diagram of our inventive technique for securely watermarking software;

FIG. 4 depicts a detailed block diagram of our inventive technique shown in FIG. 3;

FIG. 5 depicts table 500 of illustrative portions of a code image that can be applied as separate

inputs to our present invention for watermarking that entire code image;

5 FIG. 6 depicts illustrative edge and associated routine insertion in cluster flow graphs G' and H' that can arise from use of our present inventive technique, and resulting execution (control) flow among various routines, one of which having been inserted;

10 FIG. 7 depicts a high-level block diagram of computer system 700, illustratively a personal computer (PC), that can be used to implement our present invention;

15 FIG. 8 depicts a high-level flowchart of watermarked code generation procedure 800 that is executed by computer system 700 to implement our present invention;

20 FIG. 9 depicts correct alignment of the drawing sheets for FIGs. 9A and 9B;

25 FIG. 9A and 9B collectively depict a high-level flowchart of edge insertion procedure 900 that is executed as part of procedure 800 shown in FIG. 8; and

 FIG. 10 depicts a flowchart of node selection procedure 1000 which is executed as part of procedure 900 shown in FIG. 9.

To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures.

5

DETAILED DESCRIPTION

After considering the following description, those skilled in the art will clearly realize that the teachings of our present invention can be utilized in a wide range of applications for generating and incorporating a pre-defined watermark into nearly any type of computer executable code, i.e., a non-marked program (or a portion of it), or for that matter into any software object, such as an audio file, containing, e.g., music, or a video file, containing, e.g., a movie, that is to be watermarked. Through this technique, substantial tamper resistance can be imparted to the marked object itself; and other executable code, such as a pre-defined executable security code or a pre-defined executable, can be tightly integrated into the non-marked object in a manner which assures that a resulting watermarked object, which contains both the non-marked object and the watermark, is itself substantially tamper-resistant.

10

15

20

25

To clearly elucidate the invention and facilitate reader understanding, by way of background, we will first very briefly describe well known, though separate concepts of watermarking and program flow analysis. Thereafter, we will discuss our present

30

00555694.031400

invention in the context of its use for imparting a watermark, specifically implemented through pre-defined executable code, into a non-marked application program intended for subsequent execution on a personal computer (PC). During the course of this discussion, we will address the overall constituents of the watermark and its variations.

A. Overview

1. Watermarking

FIG. 1 depicts a simplified high-level block diagram of conventional watermarking process 100. In essence, an object, O, to be protected, whether it be a printed image, a document, a piece of paper currency or some other such item, is applied, as symbolized by line 105, to marking process 110 situated at an originating location. This process creates a watermark and embeds it in the object to create a watermarked object, O'. The watermarked object is then eventually transported through insecure channel 115, whether it be, e.g., transit through a public carrier or, as in the case of currency, public distribution, to a destination location. In a typical situation, a malicious attacker may intercept watermarked object O', as it travels through channel 115, and try to remove the watermark from O' by subjecting O' to his or her own transformations. Thus, an object, O" ultimately received at a destination location, may not be exactly the same, as what it was, i.e., O', when sent. At the destination, the watermarked

object O" (if a third party attempted to modify the object) or O' (if the object was not modified) is subjected to watermark recovery process 130 which attempts to recover the watermark from the object and, based on a result of the recovery process, indicates, as symbolized by output line 135, whether the watermark is present or not in the relevant object. This indication can be used to signify whether the received watermark object, then situated, as symbolized by line 120, at the destination is legitimate or not. Since the legitimacy of the document is directly governed by the security of the watermark, the watermark itself must be as difficult as possible for a third party to copy or alter.

2. Program flow analysis

FIG. 2 depicts a simplified and illustrative flow graph produced through conventional program flow analysis.

In essence and to the extent relevant, flow analysis, as symbolized by arrow 215, transforms executable code, such as program listing 210, into a connection graph, such as graph 220. Each node within the graph represents a separate instruction in the code, with typically a first such instruction encountered in the code represented by an uppermost ("root") node. Here, separate instructions in listing 210 are graphically represented by nodes 225 having separate illustrative nodes 225₁, 225₂, 225₃, 225₄, 225₅, 225₆, 225₇, 225₈, 225₉, 225₁₀ and 225₁₁. Execution flow from one

instruction to another is indicated by a line (commonly and hereafter referred to as an "edge" or link) that connects the two nodes representing these particular instructions. As shown, instructions represented by nodes 225₁ and 225₂ are connected by edge 230₁ thereby indicating that execution flow transits between these nodes, e.g., from the instruction represented by node 225₁ to that represented by node 225₂. Similarly, nodes 225₂ through 225₁₁ are collectively connected by edges 230₂ through 230₁₃. The specific nodes and their interconnecting edges, which form a collective pattern, graphically depict the execution flow among the illustrative instructions in associated executable code, here represented by nodes 225₁ through 225₁₁ within program listing 210.

Illustrative graph 220 is quite small and greatly simplified. In actuality, with relatively large application programs, corresponding flow graphs can become extremely large, with a substantial number of nodes and edges, and graphically depict highly complex execution flows.

Advantageously, the present invention is independent of any particular technique through which flow analysis is effectuated and, in all likelihood, will properly function with any one of a wide variety of techniques provided that only one such technique is consistently used throughout any single implementation of the present invention. Hence, for brevity, we will omit all details of how flow analysis is actually performed to

yield a corresponding flow graph for associated executable code.

In a similar way to the program control flow graph, one can assign a data flow graph to given code. Clearly, our invention can function with any type of flow graph assigned to such code, regardless of whether that graph is associated with data and/or control flow.

B. Inventive software watermarking technique

1. Process overview

FIG. 3 depicts a simplified high-level block diagram of our inventive technique for securely watermarking software.

In essence, our present invention produces a watermark, illustratively containing a large number of separate executable routines, that is tightly integrated into a flow pattern of non-marked executable code, e.g., an application program, through randomly establishing additional execution (control) flows in the executable code and inserting a selected one of the routines along each such flow. Our invention can generate and embed such a watermark in substantially any executable code, such as an application program, an operating system (O/S) or a software module. For purposes of illustration, we will assume, as noted above, this code is an application program for execution on a PC.

Specifically, as shown, unmarked executable code, X, that forms, e.g., an application program that is to be watermarked, is applied to watermark code generation process 300. Process 300 first generates, through conventional flow analysis tool 310 corresponding flow graphs for this code. In actuality, this process, as implemented, in a PC (or other computer system) will generate a representation of each such graph. For simplicity, the term "graph", as used herein, is now defined to encompass not only the graph itself but also all digital and other electronic representations thereof. As will be described below in conjunction with FIGs. 4 and 5, code X can contain not only the non-marked application code itself but also other executable code, such as pre-defined security code, that is to be spliced into the non-marked application code.

Once the flow graphs are generated, those graphs are applied to watermark code generator 320. Generator 320 randomly selects pairs of nodes in both of the flow graphs, inserts an edge (i.e., to establish an execution flow) between the nodes in each pair and also inserts a pre-defined routine (both collectively viewed as an executable "procedure") (r_i) from library 330 (having routines 330₁, 330₂, ... collectively labeled as r) at an appropriate one of the nodes in each such pair; thus, ultimately forming a single combined flow graph produced by generator 320. By this we mean the following. Code segments associated with routine r_i are added to the original program X so that the behaviour of the program in terms of its input and output as well as

functionality are almost the same as the original program. Further, these segments are designed to appear as if they are a normal part of original code X so that these segments are not readily apparent, unless the
5 attacker expends considerable effort to understand the code and its behaviour. An illustrative implementation could be that, e.g., routine r_{i+1} is called by some code to alter data or variables (or code) in some segment and later after some execution routine r_{i+2} is called to
10 reverse this change. Without the reversal, the watermarked program would fail to operate normally. The flow graphs of the code for the inserted routine may encode additional information which itself may also be a part of the watermark. The information in the watermark
15 may also be in various components of H' , which is discussed below.

Now, we will describe an exemplary
implementation of how one may embed information in a flow
20 graph that has a distinguished start node and where the nodes are canonically labeled using some scheme (e.g., by integers, short binary strings or elements from a finite field). From the distinguished node, one may traverse the graph using any one of many conventional traversal
25 schemes which visit all the nodes in a predetermined sequence. Illustrative ways of doing so include depth first search and breadth first search. Here, we may modify nodal ordering using a pseudo-randomized strategy. For example, where the traversal is at a particular node
30 and when deciding which node to visit next, some possible candidates may be pseudo randomly selected for omission.

This provides some resilience inasmuch as an attacker changing the code and the associated flow graph may not know which candidates are omitted and included. Further, the non-omitted candidates may be pseudo randomly ordered and then visited in that order. Any such traversal will yield a sequence (order) of nodes visited, and their nodal labels in that order will yield a parameter, i.e., a specific code (so as not to confuse the reader, the term "code" here is being used in a cryptographical sense and not referring to software code) that represents the graph. We may use more than one traversal. The parameter is likely to be distinct for each graph and its canonical labeling. To further enhance the resilience, an error correction procedure may be applied over this parameter, for example, using majority logic decoding procedures or algebraic coding procedures. In the latter case, the graphs and the labeling, and the pseudo-random choices may need to be pre-configured so that the error correction may be applied meaningfully. Then, the output of a decoding process (or a decryption process using a cipher that employs a key derived from the watermark key) would be a message, decrypted using a secret key that is embedded in the graph. The distinguished node may not be easily identifiable by an attacker who does not know the secret, whereas given the secret key one can readily identify the distinguished node. In particular, one method to accomplish this would be that while in the detection phase one tries each possible candidate node and tries to find an associated code of a subgraph of certain fixed size that contains that node as a root. If a resulting derived parameter appears meaningless or an

attempt to recover the parameter fails for some reason,
one concludes there are no watermarks; else, we conclude
that there is a watermark. Examples of such parameters
(cryptographic codes), which are well known in coding
theory are described in, e.g., J.H Van Lint, Introduction
5 to Coding Theory, (© 1998, Springer Verlag).

Furthermore, any pseudo-random number generator
needs a secret random key which may be derived from a key
10 used for watermarking using any one of many conventional
ways. For example, to derive numeric key k_1 from secret
key S , we can use $k_1 = \text{HASH}(S, 1)$, where HASH may be a
collision resistant hash function such as MD5 or SHA1.
For a description of these functions, the reader is
15 referred to A. J. Menezes et al, Handbook of Applied
Cryptography, (© 1997; CRC Press, LLC).

The watermark is collectively defined by the
routines and edges (the latter collectively implementing
20 a pattern of associated execution flows) that have been
added to the unmarked code and, where used, any
cryptographic parameter that is detected from
corresponding graphs.

Thereafter, generator 320, using the unmarked
code and code for routines r , assembles the marked
executable code, for the program, from that represented
by the combined flow graph. The resulting marked
executable code is watermarked code X' , here being a
30 watermarked application program.

Since the flow pattern of the watermark is highly intertwined with (tightly spliced into) the flow pattern of the non-marked code, the watermark is effectively impossible to either remove from the watermarked code and/or circumvent. Furthermore, since the code for the routines themselves is added in such a manner that the flow pattern of resulting watermarked code is not substantially different from that of the non-marked code, the watermark is also extremely difficult for the third party adversary to discern using, e.g., standard flow analysis tools.

To enhance tamper-resistance of the watermarked code, each routine, that constitutes a portion of the watermark, can provide, as discussed above, a pre-defined function such that if that routine were to be removed from the marked application code by, e.g., a third party adversary, then that application will prematurely terminate its execution.

Each of these routines is pre-defined, as noted above. They are usually quite compact in size (i.e., contain a relatively small number of instructions) and executes a pre-defined "small", often self-contained operation (task), such as, e.g., computing a cryptographic key for use in printing or decoding a variable, decrypting a ciphered variable, encrypting a variable or data value prior to either its use or storage, and/or validating either a stored parameter or an input value. Each of these operations is chosen so as not to require much processing time; thus, not noticeably

degrading execution of the application program itself. Multiple routines could be inserted to provide a combined functionality. In that regard, one such routine can be inserted in the program flow to encrypt a plaintext data value into an enciphered value, and another routine could be inserted into that flow at an appropriate location to decrypt the enciphered value immediately before the application program needs to use the plaintext data value. Collectively, the routines are preferably such that, for proper execution of the marked program, they must all be executed and, to a certain extent, in a given sequence. In that regard, if any one or more of these routines is removed from the watermarked application program, such as by a third party adversary, then that program will gracefully terminate its execution.

To further frustrate its detection, the code for all the inserted routines is collectively scattered approximately uniformly throughout the watermarked application program as generator 320 constructs that program from the combined flow graph. In this manner, the routines will not be centralized in any one portion of the watermarked program. Furthermore, each of these routines is written with conventional standard code "obfuscation" techniques to further camouflage their functionality.

2. Detailed process depiction

With the above in mind, FIG. 4 depicts a detailed block diagram of our inventive technique, with

FIG. 5 depicting table 500 of illustrative code portions that can be applied as separate inputs to our technique. To simplify understanding, the reader should simultaneously refer to both of these figures throughout the following discussion.

Non-marked executable code, e.g., an application program, that is to be watermarked through process 300 is split into two portions and separately applied as input code portions G and H on input lines 403 and 407, respectively. These portions, as shown in table 500, can illustratively comprise two identical images of the same non-marked code, two different halves of the same non-marked code, and a complete image of that code and the security code. Since the partitioning of the non-marked code is not critical, other fractional partitioning of the non-marked code can be used instead, such as illustratively 1/4 and 3/4 of the image for portions G and H, respectively, and with or without inclusion of the security code. The security code can also contain specific "watermark" code, e.g., executable code distinct from the non-marked application program, which, as its primary, if not sole, function, forms a portion of the watermark (or, as described below, provides other functionality totally apart from that of the application program). The security code can also be duplicated or fractionally partitioned, as desired, across both code portions G and H.

Within process 300, code portions G and H are applied to flow analyzer 310 and specifically subjected

to corresponding flow analysis operations 413 and 417, to
yield corresponding flow graphs \tilde{G} and \tilde{H} , respectively.
Both flow analysis operations employ the same
conventional flow analysis technique. Both of these flow
5 graphs are then routed to partitioner 420 where flow
graphs \tilde{G} and \tilde{H} are kpartitioned, via corresponding
kpartitioning operations 423 and 427, to yield cluster
flow graphs G' and H' , respectively. A k-partitioning
algorithm decomposes its input into k pieces where each
10 piece has approximately a same number of nodes, and a
number of edges between the pairs of pieces is as small
as possible and approximately the same. It can be
heuristically assumed that a graph over the partitioning
has a fairly rich structure, for example, as a randomly
15 generated graph having a same number of nodes and edges.
In fact, any partitioning of the original graph will
suffice for use with our invention, with that described
herein being only one such example. Both of the
kpartitioning operations are identical and can be
20 implemented by any one of a wide variety of conventional
partitioning techniques. Each of the resulting cluster
flow graphs contains k clusters (each cluster being a
"partition") of nodes where k is a pre-defined "secret"
value (illustratively 1000 for a large application
25 program having tens of Mbytes of instructions though the
exact value of k is not critical) provided, as symbolized
by line 421, as input to process 300. Partitioning the
flow graphs in this fashion renders the resulting cluster
flow graphs more manageable over flow graphs \tilde{G} and \tilde{H}
30 while increasing complexity of the cluster graphs over

that of the nonpartitioned flow graphs. Generally, cluster size (inverse of k) is selected such that interaction between different clusters, during program execution as shown by the corresponding flow graphs, is relatively low. An identical cluster size is used by partition operations 423 and 427.

Resulting cluster flow graphs G' and H' are then routed, as symbolized by lines 433 and 437, as input to edge definition operation 440. This operation establishes edges between nodes in different clusters in graphs G' and H' , and, where desired, different clusters solely in graph G' and solely in graph H' . To so do, operation 440 randomly selects, based on secret input parameter λ (where $0 \leq \lambda \leq 1$), M pairs of nodes, collectively spanning the desired clusters, and inserts an edge linking the nodes in each such pair, where M is also a "secret" value (typically a large pre-defined integer, such as illustratively 500,000 for a large application program having on the order of tens of Mbytes of instructions though the exact value of M is not critical). An illustrative algorithm for selecting these nodes, based on λ , will be described below in conjunction with FIG. 10. These nodal pairs will include those pairs having one node in a cluster in graph G' with the other node in a cluster in graph H' and, if desired, those pairs with both nodes in different clusters solely within graph G' and/or both nodes solely within different clusters within graph H' .

As each edge is established by operation 440, designations of the nodes associated with that edge are applied both to watermark generator 450 and to routine selector 470. Selector 470, given the location of each of these nodes within the execution flow of the application program, will select an appropriate routine from library 330 (r), specifically routines r_1 (also designated 330_1), ..., routine r_n (also designated 330_n) that is to be inserted along that edge and a specific node on that edge at which that routine is to be inserted. The routines so inserted alter the data or variables locally and create data dependencies that can not be easily analyzed with a data flow analysis program. For example, routine r_1 inserted in G' may place a call (and cause data dependencies using a random looking but efficiently inverted operation) to a copy of routine r_2 inserted in H' where this call is associated with an edge inserted, as described below, from G' to H' . A variable altered in r_2 may be subject to a transformation that undoes that alteration, but also in an easily inverted but random looking operation. For example, routine r_3 may compute a check-sum of a pre-defined code segment and write that sum into a variable in another segment where, e.g., routine r_5 is inserted (routine r_3 may also check it against a given value to see if the pre-defined code segment has been altered). A label (or other programming designation) of the selected routine is applied, as symbolized by line 475, by selector 470, to one input of watermark generator 455. The selected routine and the specific node at which that routine is to be inserted are such that the inserted routine will impart minimal, if

any, adverse affect on execution of the application program at that point in the program flow. In addition, since there are likely to be far fewer routines in library 330 than there are edges inserted into graphs G' and H', each routine within library 330 is selected approximately the same number of times for insertion into the non-marked application program.

Generator 450, in addition to receiving, as input, the label of the selected routine, also receives, as symbolized by line 445, for each inserted edge, identification of the nodes for that edge, and, as symbolized by lines 433 and 437, cluster flow graphs G' and H'. With this input information, generator 450 inserts the edge into cluster flow graphs G' and H' (or either one of these graphs as appropriate) to span the nodal pair and also inserts the selected routine at one of these two nodes. The insertion process is simplistically represented by combiner 455. All the edges collectively and effectively splice the two clustered flow graphs G' and H' together into a single combined flow graph J which appears at an output of generator 450 and is routed, as symbolized by line 457, to code generator 460. Executable code is then produced, by generator 460, which corresponds to that depicted in the combined flow graph J. Generator 460 is conventional in nature. In essence, generator 460, in response to the actual code for program portions G and H, appearing on respective lines 403 and 407, and the code for all the selected routines appearing on line 485, completely assembles executable code to fully implement the flow

depicted in combined flow graph J and thus produce, on
output line 490, a new, though marked, executable
version, X', of the application program. To frustrate
third party detection of any of the inserted routines,
5 generator 460 distributes the associated code for these
routines approximately uniformly throughout the marked
program such that the flow patterns for the marked and
non-marked versions of the application program do not
substantially differ from one another.

10 As noted above, the watermark is collectively
defined by the routines and edges (the edges collectively
implementing a pattern of associated execution flows)
that have been added to the unmarked code and, when used,
15 any parameter (cryptographic code) that is detected from
corresponding graphs using, e.g., the procedures (e.g.,
with or without error correction) described above. By
virtue of the manner through which the edges and routines
are inserted as well as a relatively large number of such
20 inserted edges, the resulting flow pattern for the marked
program is sufficiently complex such that the watermark
can not be readily discerned by a third party adversary.

To fully appreciate resulting execution
25 (control) flow among various inserted routines, consider
FIG. 6. This figure depicts illustrative edge and
associated routine insertion in cluster flow graphs G'
and H', and resulting control flow among various inserted
routines. For simplicity, this figure only shows one
30 inserted edge; though, in actuality, a considerable

number of edges will be inserted in a non-marked application program.

5 Graphs G' (also denoted 610) and H' (also
denoted as 620) contain illustrative clusters 613, 615
and 617; and 623, 625, 627 and 629, respectively. For
simplicity, only a few of the clusters in each graph are
explicitly shown. Each of these clusters, excluding the
code for the associated routine, represents a relatively
10 large block of code in the non-marked application
program. Specifically, clusters 623, 625 and 613
represent code blocks P_1 , P_2 and P_3 , respectively.

15 Assume that watermark generator 450 (see
FIG. 4) has inserted, as shown in FIG. 6, routine r_1 (also
denoted as 330₁) within code block P_3 , and that
routines t_1 and t_2 located in code blocks P_1 and P_2 ,
respectively, are preexisting within the non-marked
application program. Execution would normally flow in
20 the non-marked application program itself, without any
diversion, via edge 650, just from routine t_1 directly to
routine t_2 .

25 Edge definition operation 440 (see FIG. 4) has
inserted edge 630, as shown in FIG. 6, that spans a nodal
pair having separate selected nodes within G' and H' and
specifically one node, in that pair, within code
block 613 and another node, in that pair, in code
block 623. Here, as is typical, executable code for
30 routine r_1 is inserted at a node, in the pair and in
block P_3 , remote from the other node, in the pair and

001120 16522560

located in block P_1 , situated within a normal execution flow of the non-marked application program. As such, execution will flow from routine t_1 in block P_1 , via edge 630 and in a direction indicated by dashed line 635, to routine r_1 in block P_3 and, once this routine has completed its execution, will return, along this edge and in a direction shown by dashed line 640, back to block P_1 . From there, execution will follow its normal flow, within the application program and specifically via edge 650 and in a direction shown by dashed line 655, to routine t_2 in block P_2 . Thus, as one can appreciate, execution flow of the non-marked application program, as modified by our inventive technique to yield the watermarked program, will be repeatedly diverted from its normal path, i.e., that associated just with the non-marked program, to execute each of the inserted routines.

To further thwart detection of any inserted routine (such as routines r_1 through r_4), executable code that implements that routine (including appropriate jump instructions) can itself be diffused (scattered in a noncontiguous fashion, i.e., in noncontiguous locations) throughout an entire corresponding code block for a cluster, or even across multiple blocks for multiple clusters, in the non-marked application program rather than contiguously in an address space associated with or appended to a single block. Such scattering need not occur on an equal rate (e.g., two instructions of the inserted routine per every 200 instructions of the non-marked application program) throughout the block.

3. Hardware

FIG. 7 depicts a block diagram of PC 700 on which our present invention can be implemented.

As shown in FIG. 7, PC 700 comprises input interfaces (I/F) 720, processor 740, communications interface 750, memory 730 and output interfaces 760, all conventionally interconnected by bus 770. Memory 730, which generally includes different modalities, including illustratively random access memory (RAM) 732 for temporary data and instruction store, diskette drive(s) 734 for exchanging information, as per user command, with floppy diskettes, and nonvolatile mass store 735 that is implemented through a hard disk, typically magnetic in nature. Mass store 735 may also contain a CDROM or other optical media reader (not specifically shown) (or writer) to read information from (and write information onto) suitable optical storage media. The mass store stores operating system (O/S) 737 and application programs 739; the latter illustratively containing watermarked code generation procedure 800 (which incorporates our inventive technique) and routine library 330 (r) (see FIG. 5). O/S 737, shown in FIG. 7, may be implemented by any conventional operating system, such as the WINDOWS NT operating system ("WINDOWS NT" is a registered trademark of Microsoft Corporation of Redmond, Washington). Given that, we will not discuss any components of O/S 737 as they are all irrelevant. Suffice it to say, that application programs 739 execute under control of the O/S.

Incoming information can arise from two illustrative external sources: network supplied information, e.g., from the Internet and/or other networked facility, through network connection 755 to communications interface 750, or from a dedicated input source, via path(es) 710, to input interfaces 720. Dedicated input can originate from a wide variety of sources, e.g., via a dedicated link or an external source. In addition, input information can also be provided by inserting a diskette containing an input file(s) (G and H) of non-marked code and, where applicable, security code portions into diskette drive 734 from which computer 700 will access and read that file(s) from the diskette. Input interfaces 720 contain appropriate circuitry to provide necessary and corresponding electrical connections required to physically connect and interface each differing dedicated source of input information to computer system 700. Under control of the operating system, application programs 739 may exchange commands and data with the external sources, via network connection 755 or path(es) 710, to transmit and receive information, to the extent needed if at all, during program execution.

Input interfaces 720 also electrically connect and interface user input device 795, such as a keyboard and mouse, to computer system 700. Display 780, such as a conventional color monitor, and printer 785, such as a conventional laser printer, are connected, via leads 763 and 767, respectively, to output interfaces 760. The output interfaces provide requisite circuitry to

electrically connect and interface the display and
printer to the computer system. As one can appreciate,
our present inventive software watermarking technique can
be used to watermark any type of code regardless of the
modalities through which PC 700 will obtain, store and/or
communicate that code.

Furthermore, since the specific hardware
components of PC 700 as well as all aspects of the
software stored within memory 735, apart from the various
software modules, as discussed below, that implement the
present invention, are conventional and wellknown, they
will not be discussed in any further detail.

4. Software

FIGS. 810 collectively depict high-level
flowcharts of salient software procedures (modules),
which execute on PC 700, for implementing our present
invention, with specifically FIG. 8 depicting a
high-level flowchart of watermarked code generation
procedure 800. This process implements the process
provided by watermarked code generation process 300 shown
in FIG. 4. For ease of understanding, the reader should
simultaneously refer to both FIGS. 4 and 8 throughout the
following discussion.

Upon entry into procedure 800, execution first
proceeds to block 810. This block, when executed, reads
input values of secret parameters k , M and λ .
Thereafter, execution proceeds to block 820 which reads

input code portion G and then performs flow analysis on that portion to yield corresponding flow graph \tilde{G} . Once this flow graph is fully produced, then execution proceeds to block 830 which reads input code portion H and then performs flow analysis on that portion to yield corresponding flow graph \tilde{H} . Blocks 820 and 830 implement flow analysis operations 413 and 417, respectively. After cluster flow graph \tilde{H} is fully produced, block 840 is executed to kpartition flow graph \tilde{G} into cluster flow graph G'. Thereafter, block 850 is executed to kpartition flow graph \tilde{H} into cluster flow graph H'. Blocks 840 and 850 implement kpartitioning operations 423 and 427, respectively. Once block 850 fully executes, execution then proceeds to block 860.

Block 860, when executed, invokes edge insertion procedure 900 (to be discussed below in conjunction with FIGs. 9A and 9B) to randomly insert M separate edges and routines collectively between clusters in graphs G' and H', and, if desired, different clusters solely within each of graphs G' and H', to yield single combined flow graph J. This procedure, which implements edge definition operation 440, provides: random selection of nodes in both and, where desired, either one of the cluster flow graphs G' and H' to form nodal pairs; insertion of edges to connect each such nodal pair; and selection of a proper routine for insertion along the edge defined by each nodal pair and selection of a

particular node in that pair at which that routine is to be inserted.

Once procedure 900 fully executes, execution proceeds to block 870, which implements code generator 460. This block, using input code portions G and H, and library 330, constructs executable code corresponding to the representation depicted in combined flow graph J. In addition, block 870 distributes, as noted above, the code for the individual inserted routines approximately uniformly throughout the executable code. Once this executable code is fully generated, execution passes to block 880 which simply provides this code as output watermarked code. Thereafter, execution exits from procedure 800.

FIG. 9A and 9B collectively depict a high-level flowchart of edge insertion procedure 900 that is executed as part of procedure 800 shown in FIG. 8; the correct alignment of the drawing sheets for FIGs. 9A and 9B is shown in FIG. 9.

As shown in FIGs. 9A and 9B, upon entry into routine 900, execution first proceeds to block 905 which, when executed, initializes contents of counter i to one. Thereafter, execution proceeds to a loop formed of blocks 910 through 960 to: randomly select nodes that are to form nodal pairs; select a proper routine for insertion at an edge defined by each such nodal pair and select a particular node in that pair at which that

routine is to be inserted; and connect the nodes in each pair together to form M separate edges.

In particular, upon entry into this loop,
5 execution first proceeds to block 910. This block, when executed, selects a specific flow graph(s) to which edge(i) is to be added such that edges (links) are added to all such graphs in accordance with corresponding pre-defined distributions. Illustratively, such
10 distributions may call for approximately the same number of edges to be inserted that span both cluster flow graphs as are collectively contained solely within graph G' and within graph H', or a different allocation of edges between and within each of graphs. The specific
15 distributions that are used, in any one instance, are not critical provided that an adequately large number of edges is inserted that span both cluster flow graphs G' and H' so as to yield a sufficiently complex flow pattern that thwarts third party detection of the complete
20 watermark. Moreover, since all edges are effectively chosen at random, then the edges can be inserted in essentially any order between both flow graphs and, where appropriate, solely within either of the flow graphs. Once the appropriate graph(s), such as both G' and H' or
25 solely G' or solely H', are selected for edge(i), then execution proceeds to decision block 915. This block appropriately routes execution based on the selected graph(s). In particular, if edge(i) is to be inserted solely within graph G', then decision block directs
30 execution, via path 916, to block 925. Block 925 selects, typically randomly, two nodes in graph G',

though with each node in a different cluster, as endpoints for edge(i). Alternatively, if edge(i) is to be inserted that spans both graphs G' and H', then decision block directs execution, via path 918, to block 930. Block 930 selects, typically randomly and illustratively through invoking node selection procedure 1000 (to be discussed below in conjunction with FIG. 10), two nodes, one in graph G' and the other in graph H' as endpoints for edge(i). Lastly, if edge(i) is to be inserted solely within graph H', then decision block directs execution, via path 922, to block 935. Block 935 selects, typically randomly, two nodes in graph H', though with each such node in a different cluster, as endpoints for edge(i). In any instance, a node can be common to two inserted edges. Thereafter, after block 925, 930 or 935 executes, execution proceeds, via corresponding path 936, 938 or 942, to block 945. Also, once appropriate nodal pairs have been selected, in the manner described above, further pairs can be selected by randomly picking a pair of nodes from a union of the graphs G' and H' to yield further edges.

Block 945 selects, for edge(i) and particularly with respect to the location of each node in this edge in the corresponding program flow, an appropriate routine from library 330 (see FIG. 3) that is to be inserted along this edge. Once this routine is selected, execution proceeds to block 950, as shown in FIGs. 9A and 9B. This block, when executed, connects the nodal pair for edge(i) together to effectuate this edge, and selects one of these two nodes in one flow graph as an insertion

point for the selected routine, such that this routine causes minimal, if any, adverse affect on execution flow. Typically, where an inserted routine is to be invoked when program execution reaches one node in a nodal pair, a call (or jump instruction) to that routine will be located at that node but with executable code for that routine being situated at the opposite node in that pair. Once block 950 fully executes, execution proceeds to decision block 955 which determines, based on current contents of counter i, whether all M edges have been defined. Specifically, if the contents have a value less than M, then decision block 955 routes execution, via NO path 957, to block 960. This latter block increments the contents of the counter by one. Thereafter, execution loops back, via path 965, to block 910 to select appropriate flow graph(s) for the next edge, and so forth. Alternatively, if the current value of counter i equals M, then all M edges have been generated; hence, decision block 955 routes execution, via YES path 959, to block 970. This latter block, when executed, provides, as output, single combined flow graph J that through the addition of edges and routines has effectively spliced together graphs G' and H'. Thereafter, execution simply exits from routine 900.

FIG. 10 depicts a flowchart of node selection procedure 1000, which is executed as part of procedure 900 shown in FIG. 9. This procedure implements an illustrative heuristic for selecting two nodes in a nodal pair that is to span graphs G' and H'.

Upon entry into routine 1000 shown in FIG. 10, execution first proceeds to block 1010. This block, when executed, randomly picks a node, U, in cluster flow graph G'. Thereafter, execution proceeds to block 1020. Here, with λ being pre-defined as equaling (a number of edges that are to transit between G' and H')/(a number of edges connected to U), block 1020 randomly selects, with a probability of 1λ , a node, Y, in graph H'. Then, block 1030 randomly chooses, with a probability of λ , a node, Z other than U, in graph G'. Once this occurs, block 1040 executes to provide, as output, the designations of two nodes Y and Z as a nodal pair. Thereafter, execution exits from procedure 1000.

Nodal pairs may be selected using any one of many different ways, apart from that shown in FIG. 10 and discussed above. Another such way involves using bipartitie graphs. In particular, we note here that the graphs G' and H' may have clusters which may be unequal in number. From graph G', pick a pre-determined number of nodes, for each cluster, which need not be all equal. Let the number of nodes be L. Then, similarly pick L' (often a pre-determined number, but possibly given by a probabilistic distribution) nodes from graph H'. Now, construct a random bipartitie graph with the L and L' nodes as independent sets and with an average degree of each node being small and pre-determined. Then, add edges among the nodes in both of the independent sets using any standard method of random graph generation, once again keeping the average degree small and, for example,

comparable to the average degree of the nodes in the original graphs G' and H' .

The number of edges that connect to U is implied by the process above. Generally speaking, the total number of edges to be added is selected based, generally qualitatively, on how much tamper-resistance is desired while not making code growth too large or a performance penalty too significant. Thus, one must avoid inserting edges into busy sections of unmarked code (e.g., inner body of a loop). Fortunately, such "busy" sections can be readily detected using standard software profiling techniques.

Though we described our invention in terms of watermarking software intended for use with a PC, our invention is clearly not so limited in its application. In that regard, the present invention can be used to impart tamper-resistance to any program code designed to execute on nearly any type of computer, not just a PC. Such other types of computers illustratively include workstations, minicomputers, mainframe computers, and specialized computers such as industrial controllers and telephony switching system computers.

The various secrets, e.g., the values k , M and λ , as described above, for detecting the watermark may be contained in a secure processor that has secure memory (whose contents can not be altered or inferred) and is capable of secure execution (i.e., it executes given code, in such a way that its execution cannot be

altered). Alternately, this processing can be implemented using O/S features in software. Another possibility is that secure memory and secure execution can be simulated in software such that the protected object will be self-enforcing.

While we have described the inserted routines as providing functionality (operations) required by the protected program and terminating that program if any one such routine were to be removed by a third party adversary, the routines are clearly not so limited in their use. In that regard, one or more of the inserted routines can provide functionality totally unrelated to, having no affect on and completely separate from that provided by the protected program. Such functionality can include restricting access to the protected program or other protected object. See, e.g., that described in co-pending United States patent application "Passive and Active Software Objects Containing Borerresistant Watermarking", serial number 09/315,733, filed May 20, 1999 and assigned to the same assignee hereof, which is incorporated by reference herein.

Use of the present invention advantageously permits a standalone feature, such as a general security function, as implemented by one or more of the routines, to be easily and securely integrated into any non-marked computer program. In fact, many different functions could be added to a single non-marked program.

For example, if the watermark program is an application program designed to execute under a certain operating system(s), then one inserted routine could establish a network connection over, e.g., the Internet, and then transmit, over that connection, a hardware identifier for a user computer to a manufacturer of an operating system then executing on that computer; and then, in return, receive, from the software manufacturer and store in, e.g., and O/S registry, an install identifier containing a signed version of the machine identifier. Another such routine could subsequently access the hardware identifier and compare that identifier to the one contained in the install identifier. If both matched, program execution would proceed unaffected under the O/S. Alternatively, if a mismatch were detected, such as would arise if an image of the O/S were to have been copied to and is now executing on another computer (which would inherently contain a different hardware identifier), then the routine could inform the O/S appropriately and display a warning message and instruct the O/S to cease further execution of any application programs.

Clearly, a myriad number of other standalone security functions and/or other nonsecurity features could be implemented by corresponding routine(s) inserted into non-marked executable code through our inventive watermarking technique. By binding these routines very tightly to the non-marked program and with randomized insertion into and throughout the program, these features will collectively exhibit an extremely complex control

flow. This flow complexity will significantly frustrate
a third party adversary, who uses conventional software
analysis tools, in detecting and circumventing all the
functions and features implemented by the inserted
routines.

In addition, though we have described our
inventive technique as utilizing two separate input code
portions and generating a partitioned flow graph for each
such portion, our inventive technique, with readily
apparent simplifications, will function in those
situations where only one such input code portion is
used. In this case, edges and accompanying routines
would be inserted along edges defined by randomly
selected nodal pairs spanning different (and/or even the
same) clusters in just one partitioned flow graph for
that input code portion. The result would be a watermark
that is likely to be simpler, and hence easier for an
adversary to discern and thus less optimal, than that
attained through use of two separate input code portions.
Furthermore, our inventive technique could also be used
to insert edges and routines among more than two separate
input code portions, e.g., three or more, if desired.
However, the complexity of the overall processing to
accomplish this result may be excessive for the marginal
benefits that could be gained. In that regard, we
believe that use of two separate input code portions will
suffice to create a watermark that is sufficiently immune
to adversarial detection and circumvention through use of
conventional flow analysis techniques.

Moreover, though the flow graphs are described as showing control/data flow on an instruction-based level for an unmarked program and hence insertion of executable routines at an "instructional" level, where security restraints can be considerably loosened, such flow graphs can be based on control/data flow that occurs at a higher level, such as on a block-by-block basis, with each node representing a routine or other execution block formed of a group of instructions, with insertion of such routines occurring at that level.

Furthermore, though we have described our invention in the context of software-implemented objects, whether in the specific context, as described above, of executable software code or as noted above, a data object, such as, e.g., an audio or video file, these objects can also be hardware-related. In that regard, a design of an integrated circuit, such as a dedicated logic chip or other digital processing device that embodies a predefined operational flow, whether it be in a parallel, pipelined or sequential fashion, in which resulting operations and accompanying control/data flow can be graphically depicted as nodes with interconnection edges, can be viewed as such an object and hence encompassed within this term. Through use of a suitable set of library routines, which, in a similar fashion to library of routine(s) discussed above, implement added functionalities and randomized edge addition procedures, such as those discussed above, a resulting graph will be generated such that, from a standpoint of input/output functionality which the chip or device is intended to

5

10

We claim:

1. Apparatus for forming an identifier for an input object and for securely marking the input object with the identifier so as to yield a marked object, the apparatus comprising:

a processor; and

a memory having computer executable instructions stored therein; and

wherein the processor, in response to the stored executable instructions:

generates a flow representation for the input object, the representation having a plurality of nodes, said nodes representing predefined first operations performed by the input object, and connections among the nodes signifying associated flow among the predefined first operations performed by the input object;

randomly selects first and second nodes from the plurality of nodes in the representation so as to form a pre-defined number of nodal pairs, each of said pairs having one of the first nodes and a corresponding one of the second nodes; and

for each of the nodal pairs, establishes flow between the first and second nodes in said each nodal pair and inserts, in the flow so established, a selected one of a plurality of different pre-defined second operations so as to collectively define the marked object, whereby the marked object implements the predefined first operations and a plurality of selected ones of the predefined second operations, each of which has been randomly spliced into flow of the input object,

2025 RELEASE UNDER E.O. 14176

30 wherein the identifier collectively comprises all the
31 different ones of the plurality of predefined second
32 operations, and the associated execution flow associated
33 therewith and involving the nodal pairs.

1 2. The apparatus in claim 1 wherein the input object
2 comprises a software object.

1 3. The apparatus in claim 2 wherein the software object
2 comprises, input executable code, at least one
3 instruction in the input executable code is associated
4 with a corresponding one of the predefined first
5 operations, and executable code for a corresponding
6 executable procedure is associated with each selected one
7 of the predefined second operations.

1 4. The apparatus in claim 3 wherein the processor, in
2 response to the stored instructions:
3 inserts a pre-defined number of separate links and
4 designations for the selected ones of the procedures into
5 the flow representation so as to yield a combined flow
6 representation; and
7 converts, in response to said input executable code
8 and executable code for said selected ones of the
9 procedures, said combined flow representation into output
10 executable code, said output executable code being the
11 marked code.

1 5. The apparatus in claim 4 wherein the input
2 executable code comprises first and second portions
3 thereof and the flow representation comprises first and

second separate flow representations for the first and second portions of the input executable code, respectively.

6. The apparatus in claim 5 wherein:
the first portion of the input executable code comprises at least a pre-defined portion of a non-marked application program; and
the second portion of the input executable code comprises a remaining portion of the non-marked application program or pre-defined executable security code.

7. The apparatus in claim 6 wherein the processor, in response to the stored instructions, inserts executable code for the selected one procedure in noncontiguous locations in the input executable code.

8. The apparatus in claim 6 wherein the processor, in response to the stored instructions, selects the procedure from a pre-defined library of stored routines, wherein said procedure is one of the stored routines.

9. The apparatus in claim 8 wherein each of the inserted procedures implements, when executed, a pre-defined function such that if any of said inserted procedures is removed from the marked code, the marked code, when subsequently executed, will terminate its execution.

- 1
- 2
- 3
- 4
- 5

- 1
- 2
- 3

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

1 13. The apparatus in claim 12 wherein the processor, in
2 response to the stored instructions, inserts executable
3 code for the selected one procedure in noncontiguous
4 locations in the input executable code.

1 14. The apparatus in claim 12 wherein the processor, in
2 response to the stored instructions, selects the one
3 procedure from a pre-defined library of stored routines,
4 wherein said one procedure is one of the stored routines.

1 15. The apparatus in claim 14 wherein each of the
2 inserted procedures implements, when executed, a
3 pre-defined function such that if any of said inserted
4 procedures is removed from the marked code, the marked
5 code, when subsequently executed, will terminate its
6 execution.

1 16. The apparatus in claim 14 wherein at least one of
2 the inserted procedures implements, when executed, a
3 pre-defined function which is independent of
4 functionality provided by the non-marked application
5 program.

1 17. The apparatus in claim 12 wherein the processor, in
2 response to the stored instructions:
3 randomly selects a node, U, in the first cluster
4 flow representation;
5 randomly selects, with probability $1-\lambda$ (where λ is a
6 pre-defined value with $0 \leq \lambda \leq 1$), a node Y in the second
7 cluster flow representation;

2025 RELEASE UNDER E.O. 14176

8 randomly selects, with probability λ , a node Z,
9 other than U, in the first cluster flow representation;
10 and
11 provides designations of nodes Y and Z as the nodes
12 forming the nodal pair.

1 18. The apparatus in claim 12 wherein the processor, in
2 response to the stored instructions, randomly selects the
3 first and second nodes from different clusters solely
4 within the first cluster flow representation or from
5 different clusters solely within the second cluster flow
6 representation.

1 19. The apparatus in claim 4 wherein the processor, in
2 response to the stored instructions:
3 (a) partitions the flow representation into k-clusters
4 each so as to yield a cluster flow representation;
5 (b) randomly selects the first and second nodes in the
6 cluster flow representation so as to form a corresponding
7 one of the nodal pairs;
8 (c) inserts the designation for the selected executable
9 procedure at a first node in the nodal pair;
10 (d) repeats operations (b) and (c) a pre-defined number
11 of times so as to insert a pre-defined number of separate
12 procedures into the flow representation so as to yield
13 the combined flow representation.

1 20. The apparatus in claim 19 wherein the processor, in
2 response to the stored instructions, inserts executable
3 code for the selected one procedure in noncontiguous
4 locations in the input executable code.

2025 RELEASE UNDER E.O. 14176

1 21. The apparatus in claim 19 wherein the processor, in
2 response to the stored instructions, selects the
3 procedure from a pre-defined library of stored routines,
4 wherein said procedure is one of the stored routines.

1 22. The apparatus in claim 21 wherein each of the
2 inserted procedures implements, when executed, a
3 pre-defined function such that if any of said inserted
4 procedures is removed from the marked code, the marked
5 code, when subsequently executed, will terminate its
6 execution.

1 23. The apparatus in claim 21 wherein at least one of
2 the inserted procedures implements, when executed, a
3 pre-defined function which is independent of
4 functionality provided by the non-marked application
5 program.

1 24. The apparatus in claim 19 wherein the processor, in
2 response to the stored instructions, randomly selects the
3 first and second nodes from different clusters within the
4 cluster flow representation.

1 ~~25.~~ For use with a computer system having a processor
2 and a memory, the memory having computer executable
3 instructions stored therein, a method for forming an
4 identifier for input executable code and for securely
5 marking the input executable code with the identifier so
6 as to yield marked code, the method comprising the steps
7 of:

generating a flow representation for the input executable code, the representation having a plurality of nodes, said nodes representing instructions in the input executable code, and connections among the nodes signifying associated control flow among instructions in the executable code;

randomly selecting first and second nodes from the plurality of nodes in the representation so as to form a pre-defined number of nodal pairs, each of said pairs having one of the first nodes and a corresponding one of the second nodes; and

for each of the nodal pairs, establishing execution flow between the first and second nodes in said each nodal pair and inserting, in the execution flow so established, executable code for a selected one of a plurality of different pre-defined executable procedures so as to collectively define the marked code, whereby the marked code contains the input executable code and a plurality of different ones of the pre-defined procedures each of which has been randomly spliced into control flow of the input executable code, wherein the identifier collectively comprises the executable code, for all the different ones of the plurality of predefined procedures, and the associated execution flows associated therewith and involving the nodal pairs.

26. The method in claim 25 wherein the input object comprises a software object.

27. The method in claim 26 wherein the software object comprises input executable code, at least one instruction

in the input executable code is associated with a corresponding one of the predefined first operation, and executable code for a corresponding executable procedure is associated with each selected one of the predefined second operations.

28. The method in claim 27 wherein the establishing and inserting step comprises the steps of:

inserting a pre-defined number of separate links and designations for the selected ones of the procedures into the flow representation so as to yield a combined flow representation; and

converting, in response to said input executable code and executable code for said selected ones of the procedures, said combined flow representation into output executable code, said output executable code being the marked code.

29. The method in claim 28 wherein the input executable code comprises first and second portions thereof and the flow representation comprises first and second separate flow representations for the first and second portions of the input executable code, respectively.

30. The method in claim 29 wherein:

the first portion of the input executable code comprises at least a pre-defined portion of a non-marked application program; and

the second portion of the input executable code comprises a remaining portion of the non-marked

CONFIDENTIAL

7 application program or pre-defined executable security
8 code.

1 31. The method in claim 30 further comprising the step
2 of selecting the procedure from a pre-defined library of
3 stored routines, wherein said procedure is one of the
4 stored routines.

1 32. The method in claim 31 further comprising the step
2 of inserting executable code for the selected one
3 procedure in noncontiguous locations in the input
4 executable code.

1 33. The method in claim 31 wherein each of the inserted
2 procedures implements, when executed, a pre-defined
3 function such that if any of said inserted procedures is
4 removed from the marked code, the marked code, when
5 subsequently executed, will terminate its execution.

1 34. The method in claim 31 wherein at least one of the
2 inserted procedures implements, when executed, a
3 pre-defined function which is independent of
4 functionality provided by the non-marked application
5 program.

1 35. The method in claim 30 wherein the security code
2 provides functionality independent of any functionality
3 provided by the application program.

CONFIDENTIAL

36. The method in claim 30 further comprising the steps of:

- (a) generating first and second separate flow representations for the first and second portions of the input executable code;
- (b) partitioning each of the first and second flow representations into k-clusters each so as to yield first and second cluster flow representations, respectively (where k is a pre-defined integer);
- (c) randomly selecting the first and second nodes in the first and second cluster flow representations, respectively, so as to form a corresponding one of the nodal pairs;
- (d) inserting a designation for the selected executable procedure at a first node in the nodal pair; and
- (e) repeating operations (c) and (d) a pre-defined number of times so as to insert a pre-defined number of separate procedures into the first and second flow representations so as to yield the combined flow representation.

37. The method in claim 36 further comprising the step of inserting executable code for the selected one procedure in noncontiguous locations in the input executable code.

38. The method in claim 36 further comprising the step of selecting the procedure from a pre-defined library of stored routines, wherein said procedure is one of the stored routines.

1 40. The method in claim 38 wherein at least one of the
2 inserted procedures implements, when executed, a
3 pre-defined function which is independent of
4 functionality provided by the non-marked application
5 program.

1 42. The method in claim 36 wherein the first and second
2 nodes randomly selecting step comprises the step of
3 randomly selecting the first and second nodes from
4 different clusters solely within the first cluster flow

5 representation or from different clusters solely within
6 the second cluster flow representation.

1 43. The method in claim 28 further comprising the steps
2 of:

3 (a) partitioning the flow representation into k-clusters
4 each so as to yield a cluster flow representation;

5 (b) randomly selecting the first and second nodes in the
6 cluster flow representation so as to form a corresponding
7 one of the nodal pairs;

8 (c) inserting the designation for the selected
9 executable procedure at a first node in the nodal pair;
10 and

11 (d) repeating operations (b) and (c) a pre-defined
12 number of times so as to insert a pre-defined number of
13 separate procedures into the flow representation so as to
14 yield the combined flow representation.

1 44. The method in claim 43 further comprising the step
2 of inserting executable code for the selected one
3 procedure in noncontiguous locations in the input
4 executable code.

1 45. The method in claim 43 further comprising the step
2 of selecting the procedure from a pre-defined library of
3 stored routines, wherein said procedure is one of the
4 stored routines.

1 46. The method in claim 45 wherein each of the inserted
2 procedures implements, when executed, a pre-defined
3 function such that if any of said inserted procedures is

4 removed from the marked code, the marked code, when
5 subsequently executed, will terminate its execution.

1 47. The method in claim 45 wherein at least one of the
2 inserted procedures implements, when executed, a
3 pre-defined function which is independent of
4 functionality provided by the non-marked application
5 program.

1 48. The method in claim 43 wherein the first and second
2 nodes randomly selecting step comprises the step of
3 randomly selecting the first and second nodes from
4 different clusters within the cluster flow
5 representation.

1 49. A computer readable medium having computer
2 executable instructions stored therein for performing the
3 steps of claim 25.

1 50. Executable computer code securely marked with an
2 identifier and generated by a computer system, the system
3 having a processor and a memory, the memory having
4 computer executable instructions stored therein,
5 characterized by the code having being produced by the
6 steps, implemented by the processor in response to the
7 executable instructions, recited in claim 25.

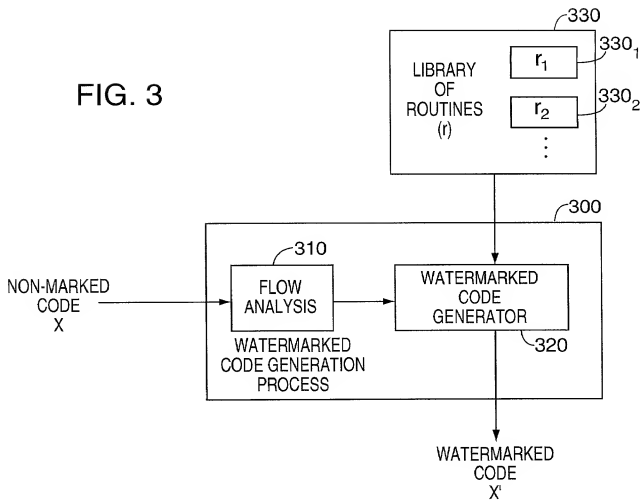
2025 RELEASE UNDER E.O. 14176

Abstract of the Disclosure

Apparatus and an accompanying method, for forming and embedding a highly tamper-resistant cryptographic identifier, i.e., a watermark, within non-marked executable code, e.g., an application program, to generate a "watermarked" version of that code. Specifically, the watermark, containing, e.g., a relatively large number of separate executable routines, is tightly integrated into a flow pattern of non-marked executable code, e.g., an application program, through randomly establishing additional control flows in the executable code and inserting a selected one of the routines along each such flow. Since the flow pattern of the watermark is highly intertwined with the flow pattern of the non-marked code, the watermark is effectively impossible to either remove from the code and/or circumvent. The routines are added in such a manner that the flow pattern of resulting watermarked code is not substantially different from that of the non-marked code, thus frustrating third party detection of the watermark using, e.g., standard flow analysis tools. To enhance tamper-resistance of the watermarked code, each such routine can provide a pre-defined function such that if that routine were to be removed from the marked code by, e.g., a third party adversary, then the marked code will prematurely terminate its execution.



FIG. 3



INPUT CODE PORTIONS	INPUTS		500
	G	H	
	COMPLETE CODE IMAGE	COMPLETE CODE IMAGE	
	1/2 CODE IMAGE	1/2 CODE IMAGE	
	COMPLETE CODE IMAGE	SECURITY CODE	
	SECURITY CODE	COMPLETE CODE IMAGE	

FIG. 5




FIG. 4




FIG. 7

FIG. 8

WATERMARKED
CODE
GENERATION
PROCEDURE

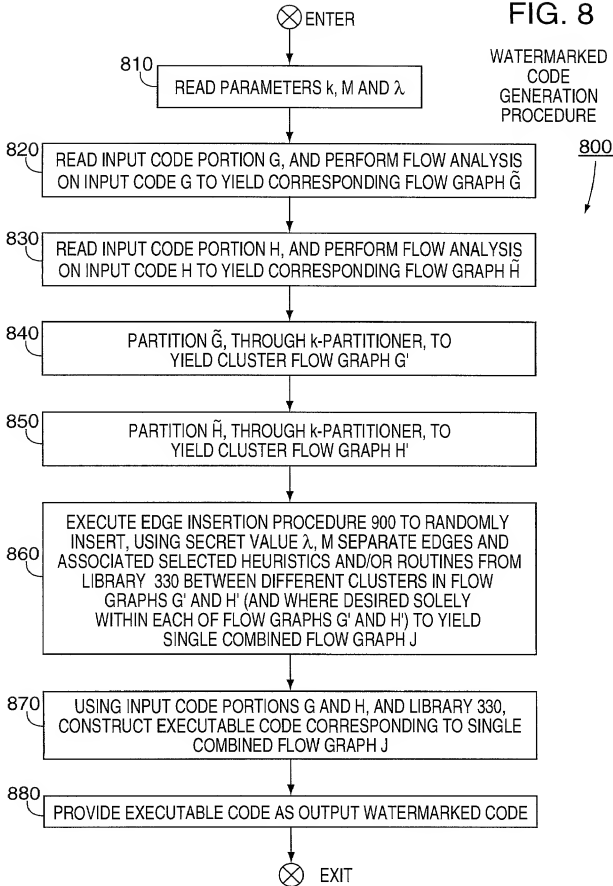


FIG. 9A

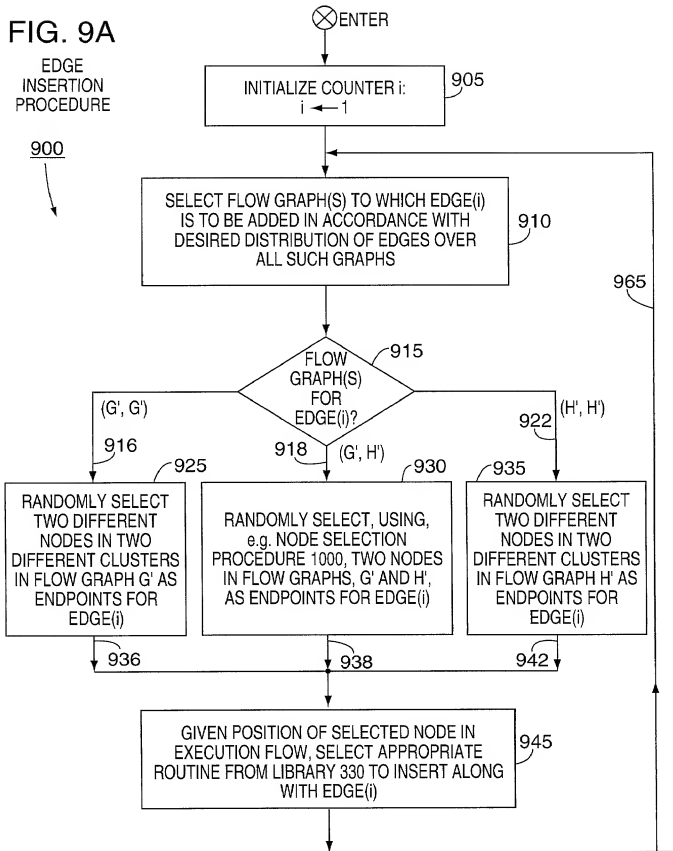
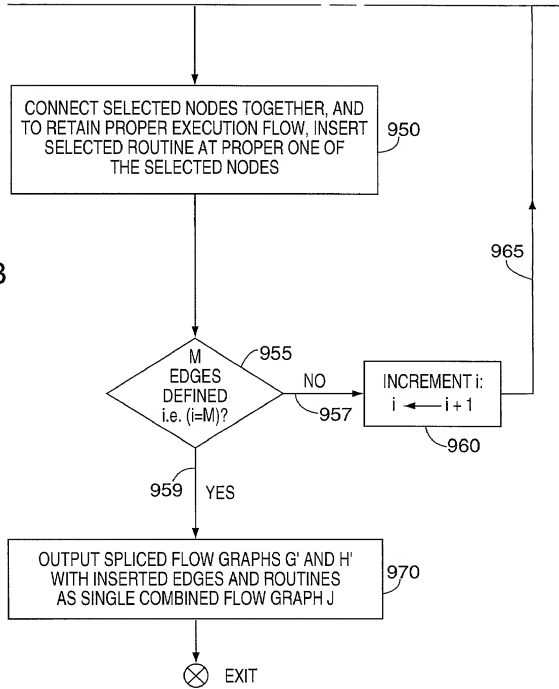
EDGE
INSERTION
PROCEDURE

FIG. 9B



9A

9B

FIG. 9

FIG. 10

NODE
SELECTION
PROCEDURE